

High-precision Monte Carlo study of the three-dimensional XY model on GPU

Ti-Yen Lan (藍棣彦),¹ Yun-Da Hsieh (謝昀達),¹ and Ying-Jer Kao (高英哲)^{1, 2, *}

¹*Center of Theoretical Science and Department of Physics,
National Taiwan University, Taipei 10607, Taiwan*

²*Center for Advanced Study in Theoretical Science,
National Taiwan University, Taipei 10607, Taiwan*

(Dated: November 6, 2012)

We perform large-scale Monte Carlo simulations of the classical XY model on a three-dimensional $L \times L \times L$ cubic lattice using the graphics processing unit (GPU). By the combination of Metropolis single-spin flip, over-relaxation and parallel-tempering methods, we simulate systems up to $L = 160$. Performing the finite-size scaling analysis, we obtain estimates of the critical exponents for the three-dimensional XY universality class: $\alpha = -0.01293(48)$ and $\nu = 0.67098(16)$. Our estimate for the correlation-length exponent ν , in contrast to previous theoretical estimates, agrees with the most recent experimental estimate $\nu_{\text{exp}} = 0.6709(1)$ at the superfluid transition of ^4He in a microgravity environment.

I. INTRODUCTION

One of the most beautiful ideas in physics is the renormalization-group (RG) theory,¹ which states that near a critical point, the nature of the phase transition can be described by a few universal properties. These universal properties depend only upon the spatial dimensionality, and the symmetry of the order parameter, regardless of the microscopic details of the system. This indicates that the phase transitions can be classified into different *universality classes*, and the asymptotic critical behaviors of each class are described by a set of critical exponents and scaling functions. Among these, the three-dimensional (3D) XY or O(2) universality class is the most extensively studied due to its relevance to the nature of the phase transitions in several physical systems, such as the λ -transition in ^4He . Experimentally, this superfluid transition permits the most accurate measurements of the critical exponents up to date in the micro-gravity environment.² The most recent value of the correlation-length exponent is $\nu_{\text{exp}} = 0.6709(1)$, which is derived from the measured value of the specific-heat exponent α via the hyperscaling relation.²

The XY universality class has been studied by various theoretical approaches: analytical field-theoretical methods^{3,4}, and numerical methods such as high-temperature (HT) expansions,⁵ and Monte Carlo (MC) simulations.⁵⁻¹⁰ MC simulations combined with the finite-size scaling (FSS) technique¹¹ have long been used to estimate the critical exponents of phase transitions. With smaller system sizes, deviation from the universal behavior due to the irrelevant scaling operators can be the source of systematic errors in the FSS analyses, and corrections to scaling become necessary. It has been proposed that the corrections to scaling can be minimized by simulating the 3D two-component ϕ^4 model on a simple cubic lattice, which belongs to the 3D XY universality class, with a proper choice of a parameter in the model.^{5,8} However, the effect of sub-leading terms can only be partially suppressed due to the limited resolution in the tuning parameter. On the other hand, by




FIG. 1. (Color online) Results for ν as a function of time. The circles show the experimental values, the upper triangle depicts the field-theoretical calculations, the squares show the Monte Carlo results, and the filled squares are the results of this work.

including sub-leading corrections in the fits, it is argued that it is possible to obtain more precise estimates of the critical exponents.⁹ The history of recent results for the correlation-length exponent ν is given in Fig. 1.

Including the sub-leading correction brings complication since it is necessary to perform a high-dimensional fit to a non-polynomial function, which might be sensitive to the numerical instability; therefore, a direct simulation of larger system sizes is desirable. Large-size simulations are crucial in developing a clear signature of criticality, and for accurate determination of the critical properties. In recent years, the advance of general purpose computing on the graphics processing units (GPUs) makes it possible to perform large-scale simulations in a massively parallel scheme.¹² In this paper, we present our MC simulations of the XY model on an $L \times L \times L$ cubic lattice up to $L = 160$ on GPU. Using data obtained from large-size systems, we are able to obtain the critical exponents with higher precision than previously achieved. This paper is

organized as follows. In Sec. II, we briefly discuss our simulation and analysis methods. Results of the simulation and a comparison with other works are presented in Sec. III. Finally, we conclude in Sec. IV.

II. MODEL AND METHOD

We simulate the classical XY model of the unit-length vectors $\mathbf{S}_i = (\cos \theta_i, \sin \theta_i)$ on an $L \times L \times L$ cubic lattice with the Hamiltonian,

$$\mathcal{H} = -J \sum_{\langle i,j \rangle} \mathbf{S}_i \cdot \mathbf{S}_j = -J \sum_{\langle i,j \rangle} \cos(\theta_i - \theta_j) \quad (1)$$

where $\langle i,j \rangle$ indicates the nearest neighbor. The periodic boundary condition is applied. In the following, we set $J = k_B = 1$.

We implement the GPU version of the Monte Carlo simulation based on the NVIDIA CUDA framework.¹³ We refer interested readers to available literature for an introduction to the details of the GPU hardware and the programming models.¹³ We implement three update schemes for the GPU version of the Monte Carlo simulation: parallel Metropolis single-spin flip, over-relaxation, and parallel-tempering methods. One Monte Carlo step (MCS) is defined as one Metropolis sweep and one over-relaxation sweep of the entire lattice, followed by one parallel-tempering exchange. To implement the parallel Metropolis and over-relaxation updates suitable for the GPU, we divide the entire lattice into blocks of $16 \times 16 \times 16 = 4096$ spins. Each block is decomposed into sub-blocks which belong to two different sub-lattices (Fig. 2). Each block is assigned to a *thread block*¹³ containing $8 \times 8 \times 8 = 512$ threads, which execute the same GPU *kernel* in parallel.¹³ Each thread is responsible for updating $2 \times 2 \times 2 = 8$ spins, with four black sites and four white sites, so that there are enough arithmetic operations to hide the latency of the global memory accesses.¹³ We apply the checkerboard decomposition algorithm to perform the Metropolis single-spin flips in parallel.^{14,15} We first update all the black spins in parallel via a GPU kernel. After all the black spins belonging to different blocks are updated, another kernel is launched to update all the white spins. Due to the special architecture of the GPU, the commonly used Mersenne-Twister (MT) random number generator (RNG) can not be efficiently implemented at the thread level. Instead, we use a faster RNG implementation specially designed for the GPU architecture, the *Warp Generator*.¹⁶ We note that although it has a smaller period of $2^{1024} - 1$ than MT ($2^{19937} - 1$), we do not find any noticeable statistical bias compared with the CPU runs using MT.

It is well established that the single-spin flip Metropolis update suffers from the critical slowing down near the critical point and one has to resort to cluster updates.^{17,18} However, implementation of the cluster update on GPU is complicated and less efficient.¹⁹ We instead implement




FIG. 2. (Color online) Checkerboard decomposition in 3D. The full lattice is decomposed into blocks containing $16 \times 16 \times 16 = 4096$ spins each. Each block is assigned to a thread block containing $8 \times 8 \times 8 = 512$ threads. Each thread manipulate $2 \times 2 \times 2 = 8$ spins, four black sites and four white sites.

the microcanonical over-relaxation update.^{6,20} The new value of the spin on site i is obtained by reflecting the spin at its local molecular field $\mathbf{H}_i = -\sum_{\langle ij \rangle} \mathbf{S}_j$,

$$\mathbf{S}'_i = -\mathbf{S}_i + 2 \frac{\mathbf{S}_i \cdot \mathbf{H}_i}{H_i^2} \mathbf{H}_i \quad (2)$$

This update maps the system from a point in the phase space to another point with exactly the same energy. After several sweeps, the system is able to explore a larger region of the phase space without being stuck in a particular local minimum, and the simulation becomes ergodic.

To better equilibrate the simulation, we also perform the parallel-tempering (PT) Monte Carlo.²¹ In the PT scheme, many replicas at different temperatures are simulated simultaneously. After a certain number of MCSs, we swap two adjacent configurations X_m, X_n at neighboring temperatures T_m, T_n with the acceptance probability of

$$W(X_m, T_m | X_n, T_n) = \min \left[1, e^{(1/T_m - 1/T_n)(E_m - E_n)} \right], \quad (3)$$

where E_n is the total energy of replica n .

We measure the following quantities during the simulation: magnetization m , Binder ratio U_L and spin stiffness ρ_s with the following estimators,

$$m = \frac{1}{N} \left\langle \sqrt{\left(\sum_j \sin \theta_j \right)^2 + \left(\sum_j \cos \theta_j \right)^2} \right\rangle, \quad (4)$$

$$U_L = 1 - \frac{\langle m^4 \rangle}{3 \langle m^2 \rangle^2}, \quad (5)$$

$$\rho_s = \frac{1}{3N} \sum_{\hat{\mu}} \left[\left\langle \sum_{\langle ij \rangle} \cos(\theta_i - \theta_j) (\hat{\epsilon}_{ij} \cdot \hat{\mu})^2 \right\rangle - \frac{1}{T} \left\langle \left(\sum_{\langle ij \rangle} \sin(\theta_i - \theta_j) (\hat{\epsilon}_{ij} \cdot \hat{\mu}) \right)^2 \right\rangle \right], \quad (6)$$

where $\hat{\mu} = \hat{x}, \hat{y}$ and \hat{z} , and \hat{e}_{ij} is the unit vector connecting nearest-neighbor sites i and j .

To reduce the amount of data transfer between the CPU and the GPU, we store all the spin configurations at different temperatures in the GPU global memory, and all updates are performed through the kernel functions on the GPU. Measurements are also performed on the GPU and the results are sent back to the CPU for binning. Simulations are carried out at 34 temperatures ranging from $T = 2.1$ to $T = 2.3$ for $L = 64, 80, 96$, and 128 and at 13 temperatures ranging from $T = 2.19$ to $T = 2.21$ for $L = 160$. The temperature set is chosen such that the acceptance rate of swaps is independent of the temperatures. After 3×10^6 MCSs for equilibrium, 1.3×10^7 measurements are made for $L = 64, 80$, and 96 and 2×10^7 measurements are made for $L = 128$ and 160. The data are blocked into several bins, each consisting of 10^5 measurements, for further analyses. Error bars are given by one standard deviation. The simulations were performed on Nvidia Tesla M2090, and took approximately 110 days of GPU time in total to accumulate the whole data set.

III. RESULTS

We perform the finite-size scaling analyses to extract the critical behaviors in the thermodynamic limit.^{11,22} The singular part of the free energy with critical exponent κ in zero-field can be describe by the scaling ansatz

$$F(t, L) = L^{\kappa/\nu} \mathcal{F}^0(tL^{1/\nu}), \quad (7)$$

where ν is the correlation-length critical exponent, $t = (T - T_c)/T_c$ is the reduced temperature, and $\mathcal{F}^0(x)$ is an universal function which is analytic as $x \rightarrow 0$.

We first use the spin stiffness ρ_s and the Binder cumulant U_L to estimate T_c . The spin stiffness ρ_s scales as

$$\rho_s = L^{-1} \mathcal{H}^0(tL^{1/\nu}) \quad (8)$$

for the 3D XY model. This indicates that $\rho_s L$ at different sizes should intersect at T_c (Fig. 3). Also, the Binder cumulant U_L for different sizes also intersect at T_c (Fig. 4). In both cases, we obtain the crossing of the curves at a consistent $T_c = 2.2019$.

Further refinement of the analysis is carried out by data collapse. We first perform data collapse on the magnetization m (Fig. 5) to obtain T_c , β and ν . Near T_c , m has the scaling form

$$m(t, L) = L^{-\beta/\nu} \mathcal{M}^0(tL^{1/\nu}), \quad (9)$$

where \mathcal{M}^0 is an universal function. We use bootstrap resampling technique²³ to decorrelate the data obtained by the PT at different temperatures. In each resampling, 1000 values are randomly chosen from the 130 (or 200) measurements and then take the average. The resampling is repeated for 130 (or 200) times to generate a new




FIG. 3. (Color online) $\rho_s L$ vs temperature for $L = 64, 80, 96, 128$ and 160. The error bars are smaller than the symbols.




FIG. 4. (Color online) U_L vs temperature for $L = 64, 80, 96, 128$ and 160. The error bars are smaller than the symbols.

data set. This data set is used to perform data collapse to obtain T_c , β and ν . Temperatures between 2.1990 and 2.2050 are used for data collapse. A fifth-order polynomial is used to fit the scaling function \mathcal{M}^0 , and the phase space of T_c , β and ν is scanned for the best collapse, where the reduced chi-square $\chi_{\text{red}} = \chi/\text{d.o.f.}$ approaches one. The procedure is repeated 100 times to estimate the error bars of T_c , β and ν (Fig. 5). Our estimates are $T_c = 2.201852(1)$, $\beta = 0.34910(12)$ and $\nu = 0.67138(11)$ with the average $\chi_{\text{red}}^2 \approx 1.2938$. We use the hyperscaling relation²⁴

$$2 - \alpha = 3\nu = 2\beta + \gamma, \quad (10)$$

to obtain the estimates for $\alpha = -0.01414(33)$ and $\gamma = 1.31594(41)$.

We also perform data collapse on q_2 defined as,

$$Q_2(t, L) = 3(1 - U_L) = \frac{\langle m^4 \rangle}{\langle m^2 \rangle^2} = q_2(tL^{1/\nu}). \quad (11)$$

TABLE I. Comparison of the critical exponents determined via various methods for three-dimensional XY universality class. The quantities with asterisk are calculated using the scaling relation $\gamma = (2 - \eta)\nu$ or the hyperscaling relation [Eq. (10)], and errors are calculated by treating variables as independent.

Method	Ref.	α	β	γ	ν
MC	this work	-0.01293(48)*	-	-	0.67098(16)
		-0.01414(33)*	0.34910(12)	1.31594(41)*	0.67138(11)
MC + IHT ^a	[10] (2006)	-0.0151(3)*	0.3486(1)*	1.3178(2)	0.6717(1)
MC	[9] (2006)	-0.0151(9)*	-	-	0.6717(3)
MC+IHT	[5] (2001)	-0.0146(8)*	0.3485(3)*	1.3177(5)	0.67155(27)
MC	[5] (2001)	-0.0148(15)*	0.3485(2)*	1.3177(10)*	0.6716(5)
ϕ^4 ^b	[4] (2001)	-0.0112(21)	0.3474(11)*	1.3164(8)	0.6704(7)
MC	[8] (1999)	-0.0169(33)*	0.349(2)*	1.3190(24)	0.6723(11)
PRG ^c	[7] (1993)	-0.014(21)*	-	1.307(14)*	0.662(7)
MC	[7] (1993)	-	-	1.324(1)	-
exp ^d	[2] (2003)	-0.0127(3)	-	-	0.6709(1)
exp	[2] (1996)	-0.01056(38)	-	-	0.67019(13)

^a improved high-temperature expansion

^b two-component ϕ^4 field theory

^c phenomenological RG

^d experiment of ^4He superfluid




FIG. 5. (Color online) Scaling of m and q_2 for $L=64$ (black), 80 (red), 96 (green), 128 (blue) and 160 (orange). The dashed lines correspond to polynomial fits to the corresponding scaling functions.

The same procedure of resampling is applied as above. Data at temperatures between 2.1990 and 2.2040 are used for data collapse. A fourth-order polynomial is used to fit q_2 . The phase space of T_c and ν is scanned to produce the best collapse. The procedure is repeated 100 times to estimate the error bars for T_c and ν . Figure 5 also shows the scaling plot of q_2 for $L = 64, 80, 96, 128$ and 160. Our estimates give $T_c = 2.2018312(6)$, $\nu = 0.67098(16)$ with average $\chi^2_{red} \approx 1.2458$. Using the hyperscaling relation Eq. (10), we obtain the estimates for $\alpha = -0.01293(48)$.

In Table I, we compare our estimates with recent results for the critical exponents of the 3D XY universality class. Our estimates of ν are consistent with each other within two standard deviations (Fig. 1), and the error bars are comparable with the experimental preci-

sion. However, they are smaller than previous theoretical estimates for ν ,^{9,10} and, contrary to previous claims, are consistent with the experimental estimate². This might be attributed to the largest system sizes ($L = 128, 160$) with high statistics that we can simulate, and a better finite-size scaling analysis can be performed.

IV. CONCLUSION

We perform large-scale Monte Carlo calculations of the 3D XY model. Implementation of efficient, highly parallel Monte Carlo update schemes on GPU enables us to perform simulations on lattices up to $L = 160$. With larger system sizes, we are able to perform finite-size scaling and obtain a five-digit accuracy of the critical exponents in a significantly less amount of computation time. With the current accuracy for ν and α in our simulations, contrary to previous theoretical studies, our results suggest that the theoretical estimates of the critical exponents for the 3D XY universality class are consistent with experimental results within two standard deviations. This suggest that the λ -transition in He^4 can indeed be accurately described by the 3D XY universality class. It would be interesting to revisit the models studied previously to further confirm results obtained in this paper.

ACKNOWLEDGMENTS

We thank A. W. Sandvik for useful discussions. This work is partially supported by NSC in Taiwan through Grant No. 100-2112-M-002-013-MY3 (Y. D. H., Y.J. K.), and by NTU Grant numbers 101R891004 (Y.J. K.).

Travel support from NCTS in Taiwan is also acknowledged.

* yjkao@phys.ntu.edu.tw

¹ K. G. Wilson, Rev. Mod. Phys. **47**, 773 (1975).

² J. A. Lipa, J. A. Nissen, D. A. Stricker, D. R. Swanson, and T. C. P. Chui, Phys. Rev. B **68**, 174518 (2003); J. A. Lipa, D. R. Swanson, J. A. Nissen, T. C. P. Chui, and U. E. Israelsson, Phys. Rev. Lett. **76**, 944 (1996).

³ J. C. L. Guillou and J. Zinn-Justin, Phys. Rev. B **21**, 3976 (1980).

⁴ H. K. Florian Jasch, J. Math. Phys. **42**, 52 (2001).

⁵ M. Campostrini, M. Hasenbusch, A. Pelissetto, P. Rossi, and E. Vicari, Phys. Rev. B **63**, 214503 (2001).

⁶ Y. H. Li and S. Teitel, Phys. Rev. B **40**, 9122 (1989).

⁷ A. P. Gottlob and M. Hasenpusch, Physica A **201**, 593 (1993).

⁸ M. Hasenbusch and T. Török, J. Phys. A: Math. Gen **32**, 6361 (1999).

⁹ E. Burovski, J. Machta, N. Prokof'ev, and B. Svistunov, Phys. Rev. B **74**, 132502 (2006).

¹⁰ M. Campostrini, M. Hasenbusch, A. Pelissetto, and E. Vicari, Phys. Rev. B **74**, 144506 (2006).

¹¹ M. E. Fisher and M. N. Barber, Phys. Rev. Lett. **28**, 1516 (1972).

¹² See <http://gpgpu.org> for articles and discussion of general-purpose computing with GPUs.

¹³ http://www.nvidia.com/object/cuda_home_new.html; D. B. Kirk and W. W. Hwu, *Programming Massively Parallel Processors* (Elsevier, Amsterdam, 2010).

¹⁴ T. Preis, P. Virnau, W. Paul, and J. J. Schneider, Journal of Computational Physics **228**, 4468 (2009).

¹⁵ M. Weigel, Computer Physics Communications **182**, 1833 (2011).

¹⁶ http://cas.ee.ic.ac.uk/people/dt10/research/rngs-gpu-warp_geo

¹⁷ R. H. Swendsen and J.-S. Wang, Phys. Rev. Lett. **58**, 86 (1987).

¹⁸ U. Wolff, Phys. Rev. Lett. **62**, 361 (1989).

¹⁹ M. Weigel, "Connected component identification and cluster update on gpu," (2011), arXiv:1105.5804v2.

²⁰ M. Creutz, Phys. Rev. D **36**, 515 (1987).

²¹ K. Hukushima and K. Nemoto, J. Phys. Soc. Jpn. **65**, 1604 (1996).

²² K. Binder, Phys. Rev. Lett. **47**, 693 (1981).

²³ C. F. J. Wu, Annals of Statistics **14**, 1261 (1986).

²⁴ M. E. Fisher, Rev. Mod. Phys. **70**, 653 (1998).